Why We Need HyperTransport in Networking Applications

John F. Wakerly
VP & CTO, Enterprise Line of Business
Cisco Systems
wakerly@cisco.com

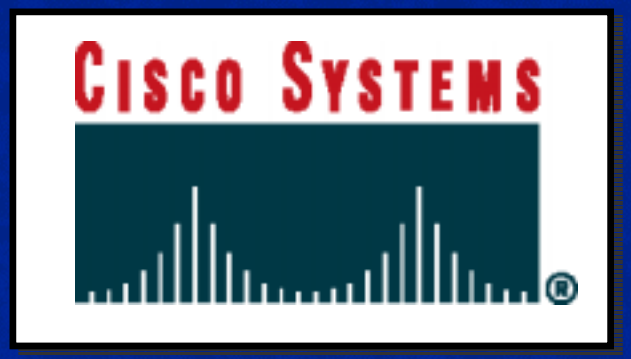

Platform Conference
Direction • Design • Perspective • Analysis

Cisco's HT Plans

- **Chip-to-chip Interface**
  - CPU to I/O devices (encryption, packet classification, etc.)
  - CPU-to-CPU (connect two SoCs back-to-back)
  - Switched applications (multiple CPUs and I/O devices)
- **Standard and Cisco-proprietary devices**

Cisco's HT Plans (2)

- **Board-to-board interface**
  - No current plans, but could be attractive if/when the need arises
  - Current products (up to OC-12 data rates) use "port adapters"
    - Standard 32-bit PCI in a Cisco-proprietary mechanical form factor
    - Strong benefits in using industry-standard chips and other infrastructure

PCI Strengths, Weaknesses

- **Strengths**
  - Well understood
  - Rich infrastructure
  - Supports general device addressing
- **Weaknesses**
  - Low bandwidth
  - Too many pins

Increasing Bandwidth / Pin

- **Point-to-point, unidirectional wiring**
  - Separate connections for transmit and receive
  - Simplified electrical termination
  - Eliminate dead states for bus turn-around
- **Differential signaling**
  - Two wires per logical signal
  - High noise immunity
  - Better tolerance to varying logic thresholds
- **Much higher-frequency operation**
- **Initial HT implementations at 800 Mb/s per pin pair, targeting over 1.2 Gb/s**

Cisco Experience with High-Speed I/O (Cat6500)

- **Cisco has shipped backplane interfaces at 1.25 Gb/s per pin pair since 1999.**
  - Interconnect from chip to connector, across backplane, through another connector and into another chip.
  - Standard FR4 PCB and backplane materials.
  - Self-clocked, rather than separate clock as in HT, but most of the experience is quite relevant.
- **Lab work indicates we can go faster, even with FR4.**

Synergy with Other I/Os

- **SPI-4 Phase 2**
  - Point-to-point streaming interface for 10 Gb/s chip-to-chip network applications (OC-192, 10GE)
  - LVDS differential signaling (2.5-V swing)
  - Source-synchronous clocking, DDR data
    - 622 Mb/s minimum data rate per pin pair
    - 16-bit data width
- **Other high-speed I/Os looking to p-to-p, differential, source-sync, DDR schemes**

Streaming vs. Addressable I/O Interfaces (1)

- **Streaming interfaces read or write an entire packet at a time**
  - Simple, but…
  - No preemption once a long packet starts.
  - Memory address controlled by DMA engine near the memory.

Streaming vs. Addressable I/O Interfaces (2)

- **Addressable interfaces provide a memory address for each transfer.**
  - Transfer size as small as a single byte.
  - Largest individual transfer size is a “block” (typically 128-256 bytes).
  - “Intelligent” peripherals have fine-grained control.
  - Peripherals can share data structures with CPUs.

High-Speed I/O in CPU-Based Network Apps

- **Traditional PC or server app:**
  - Read dataset
  - Crunch, crunch, crunch (in cache)
  - Update dataset
- **Network application:**
  - Read packet
  - Do various table lookups
  - Update packet
  - Update statistics
- **More sensitive to memory & I/O perf.**

Typical CPU-based Networking Application

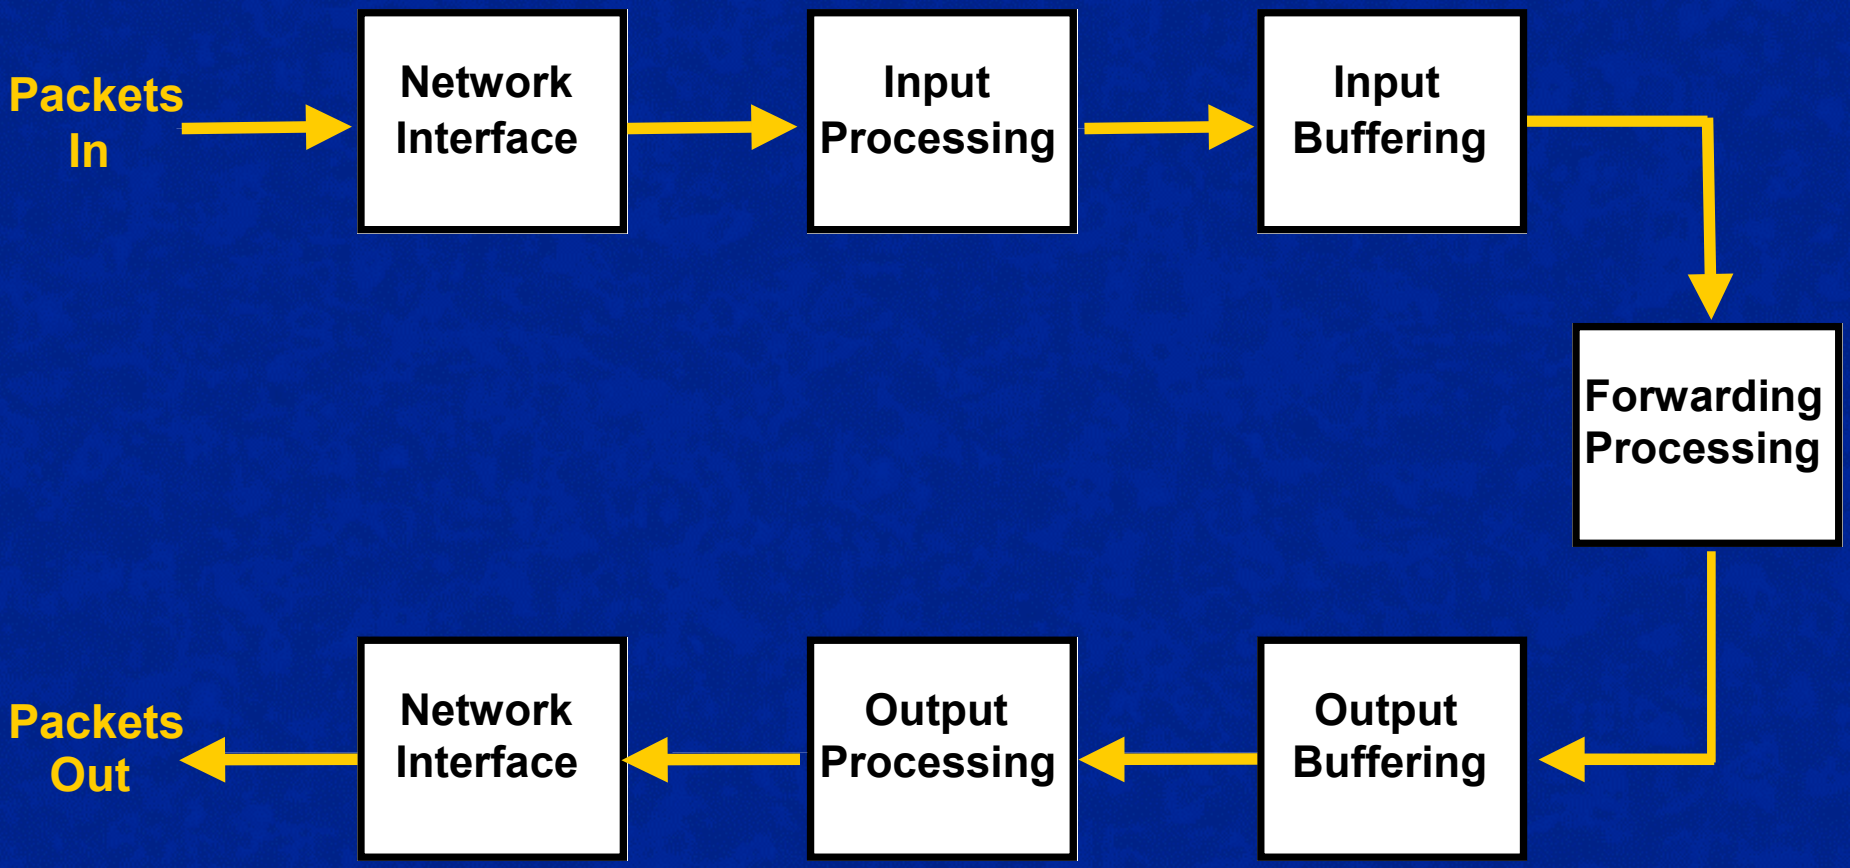

Network CPU Configuration

Note: Separate I/O and memory buses

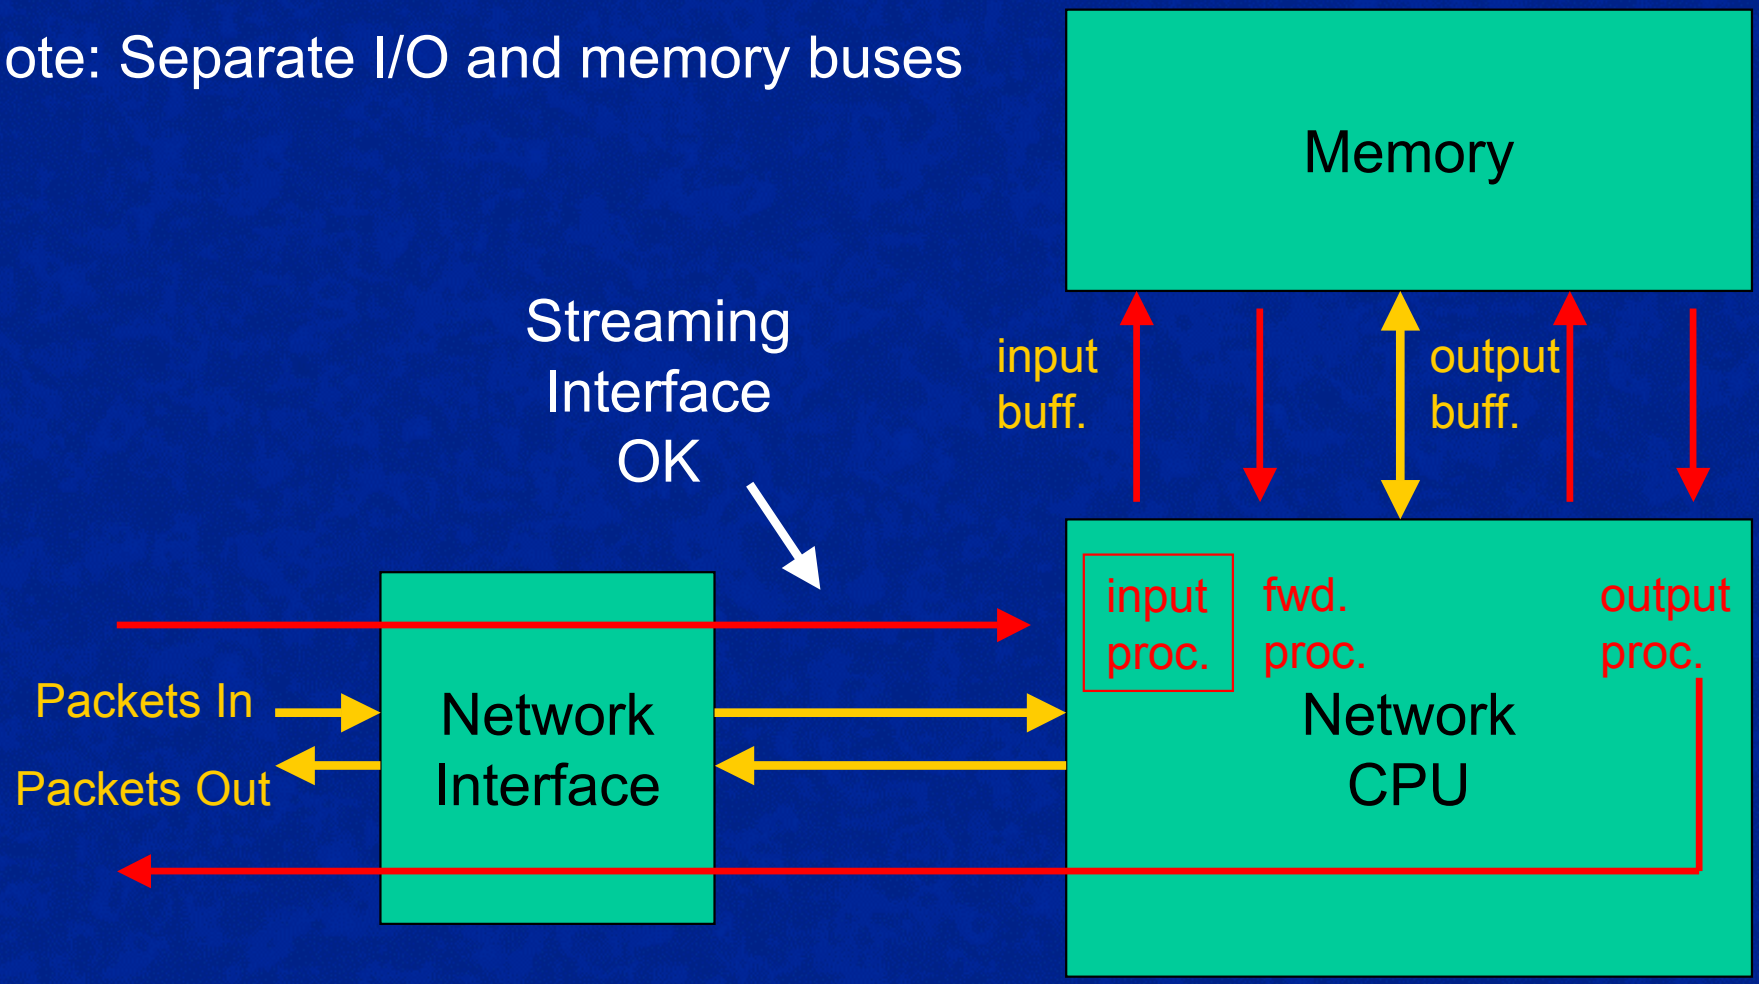

Network CPU possibilities with HT

Intelligent, high-bandwidth network interface

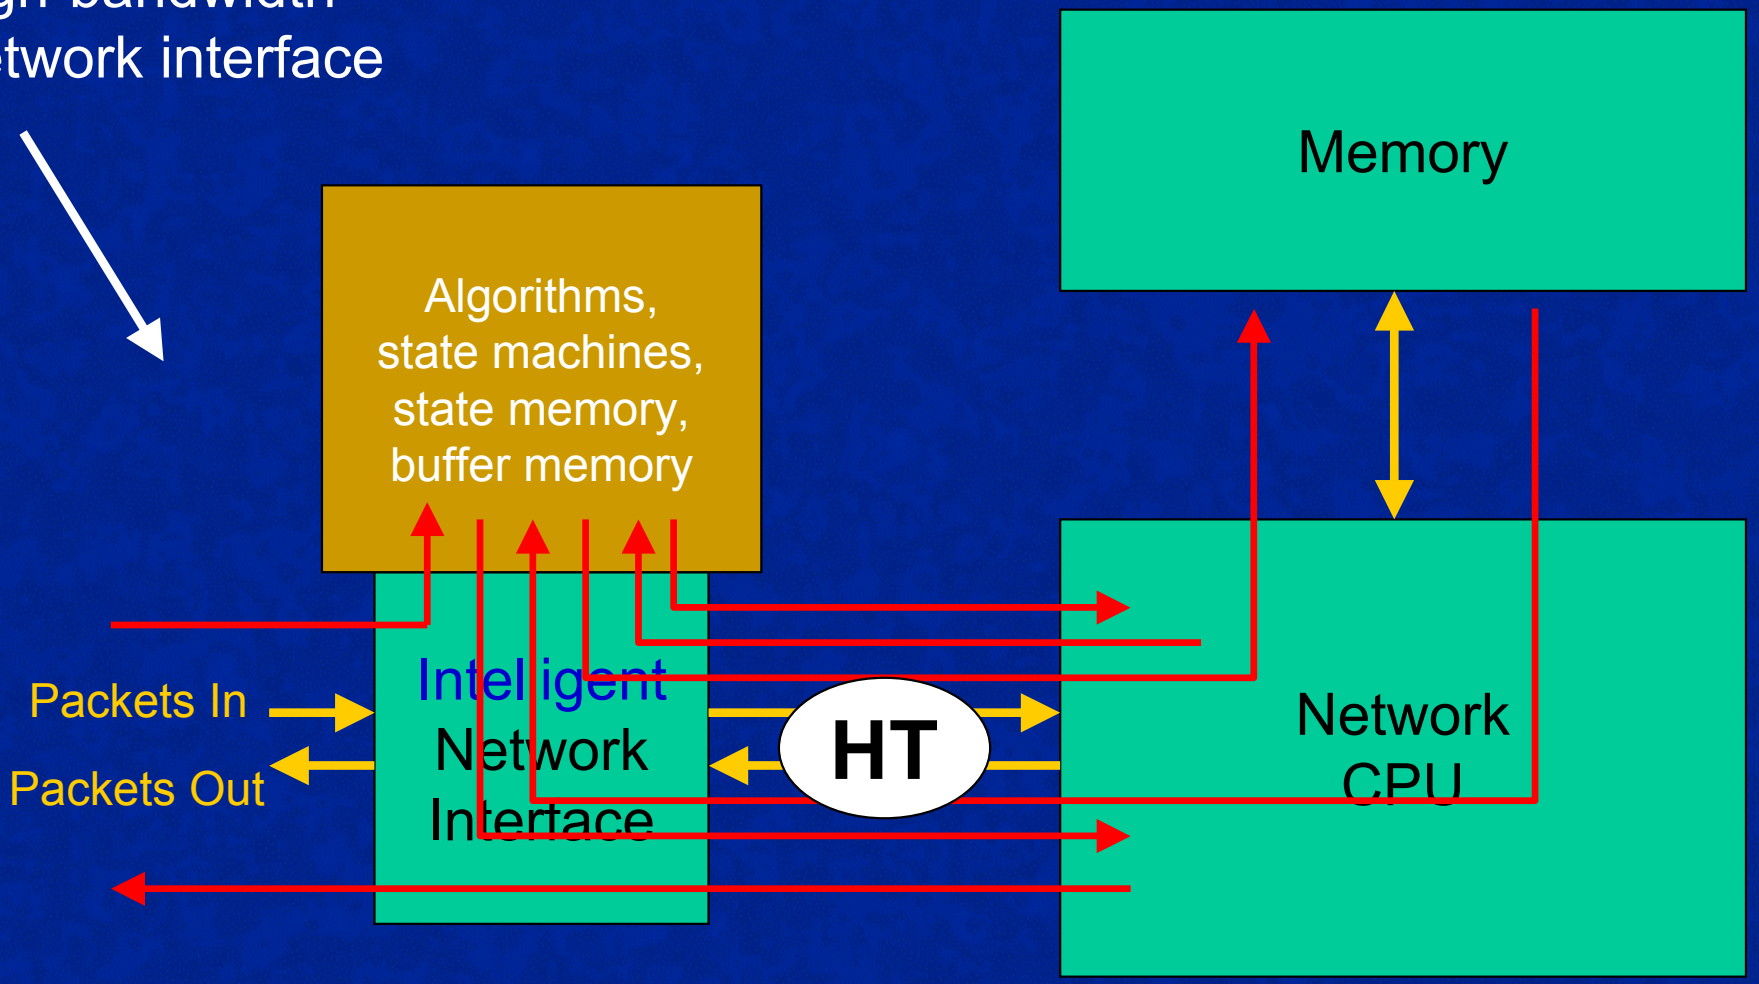

Simple networking application

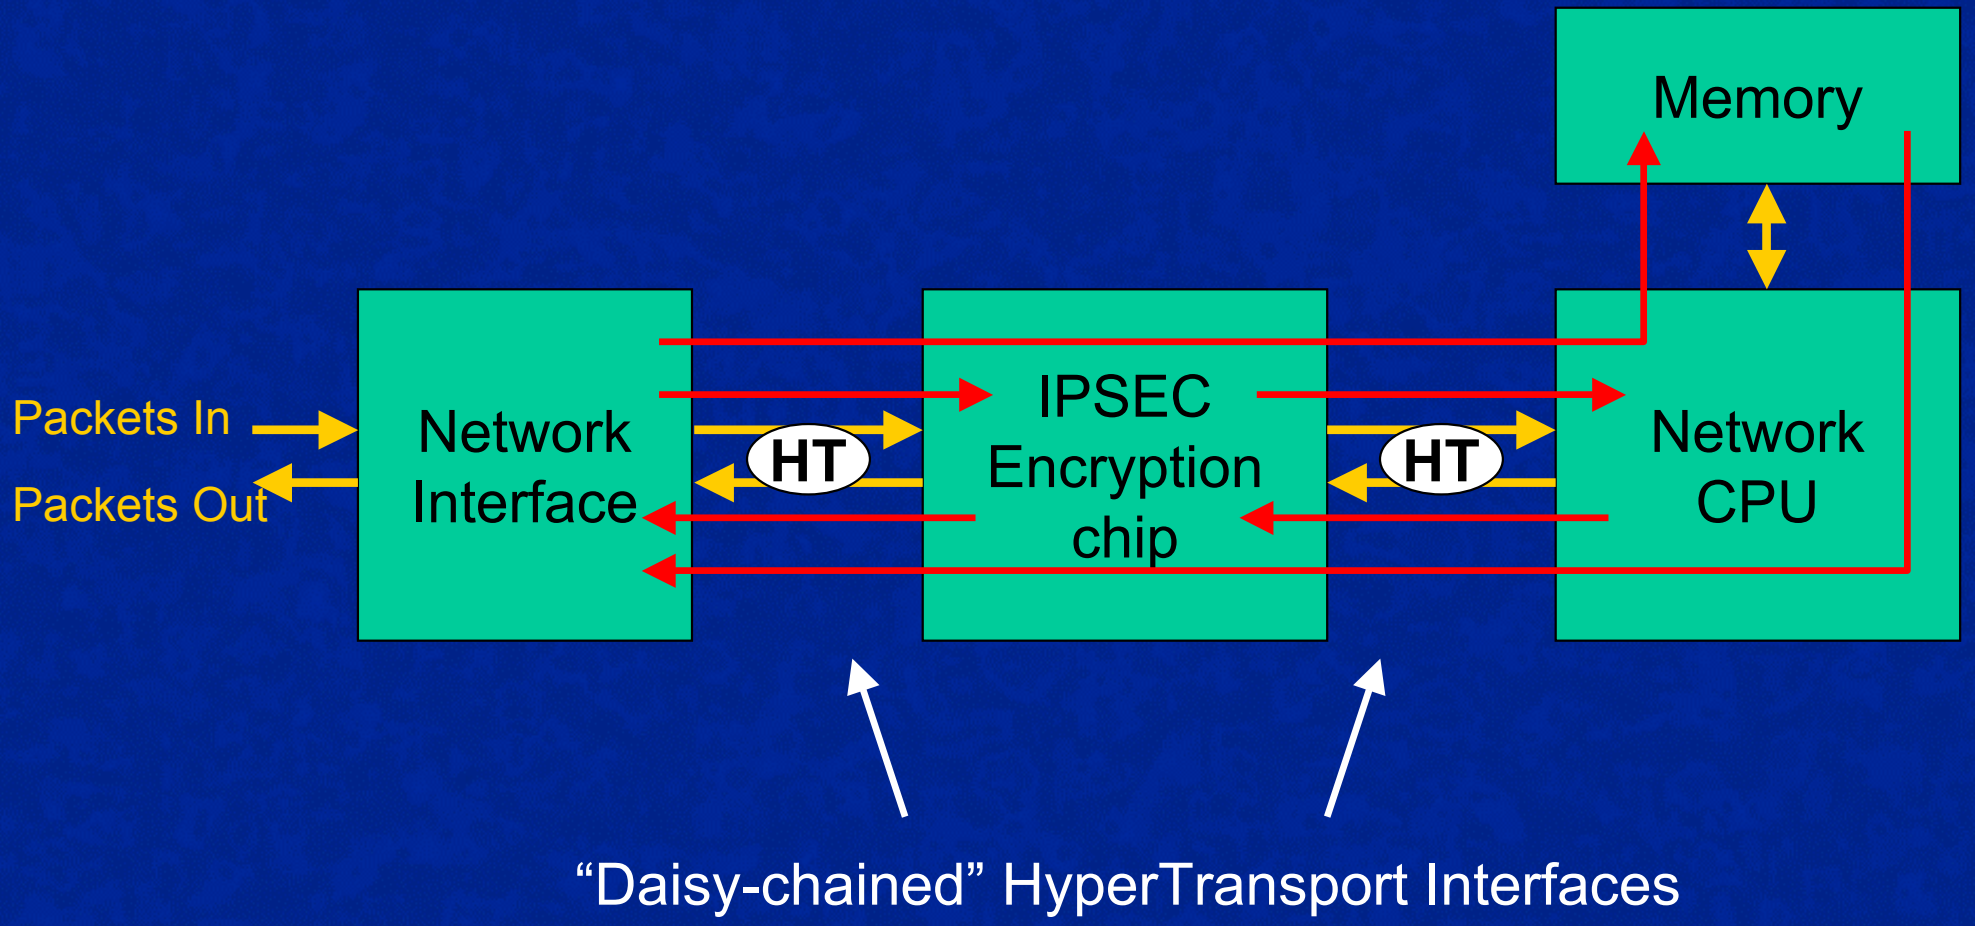

Switched networking application

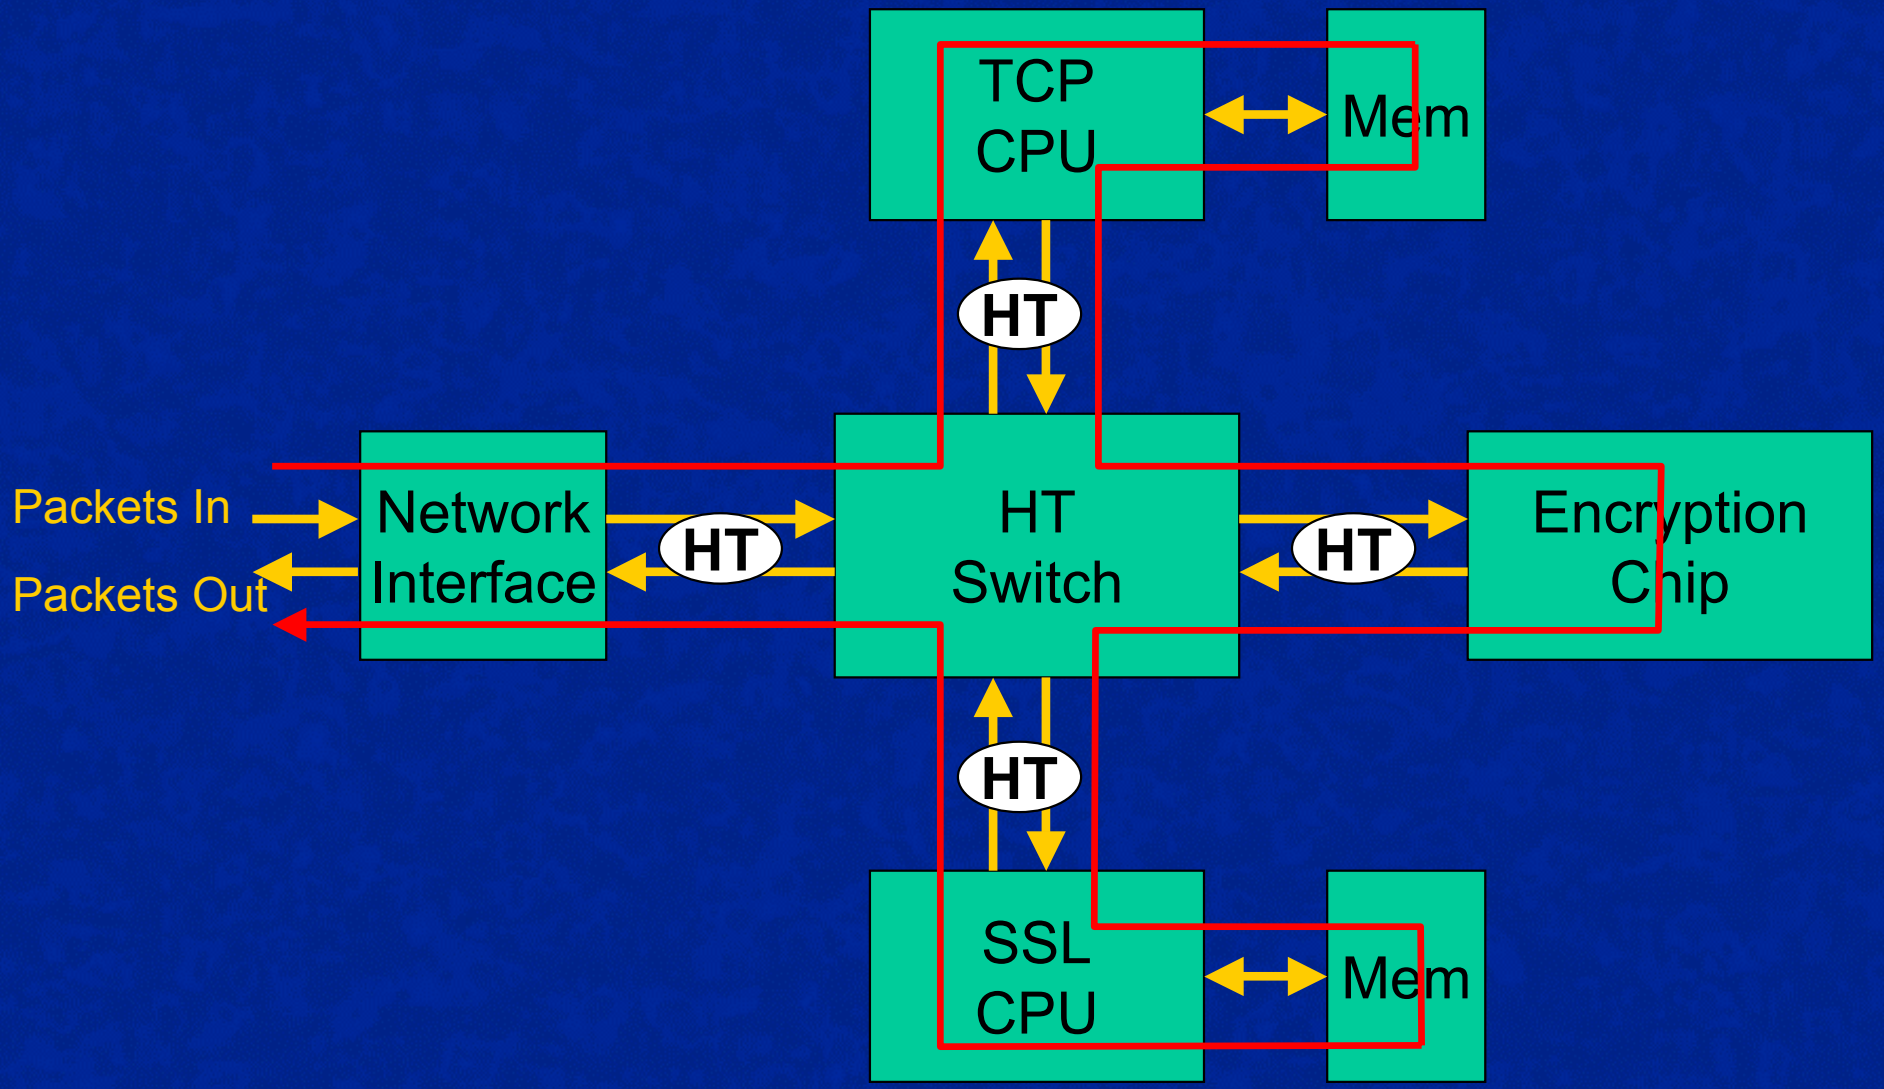

Concluding Thoughts

- **HyperTransport scalability**
  - Higher speeds, wider I/O
  - High performance “entry point”
  - Switching
- **Rich architectural possibilities**
- **Leverage – standardization across multiple industry segments**
- **Thriving collection of HT suppliers – IP, CPU, bridge, peripheral, ASIC, programmable, …..**